

EXPRESS MAIL MAILING LABEL NO. EL291391423US

PATENT

10830.0079.NPUS00

APPLICATION FOR UNITED STATES LETTERS PATENT

for

EDGE DETECTION BASED ON
VARIABLE-LENGTH CODES OF BLOCK CODED VIDEO

By

Seyfullah H. Oguz

Ugur Sezer

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to processing of compressed visual data, and in particular the processing and storage of compressed visual data for playing, transmission, or editing of an MPEG data stream.

2. Background Art

It has become common practice to compress audio/visual data in order to reduce the capacity and bandwidth requirements for storage and transmission. Some of the most popular audio/video compression techniques are defined by the MPEG family of standards. MPEG is an acronym for the Moving Picture Experts Group, which was set up by the International Standards Organization (ISO) to work on the compression of audio/visual information. MPEG provides a number of different variations (MPEG-I, MPEG-2, etc.) to suit different bandwidth and quality constraints. MPEG-2, for example, is especially suited to the storage and transmission of broadcast quality television programs.

For the video data, MPEG provides a high degree of compression (up to 200:1) by transforming 8 x 8 blocks of pixels into a set of discrete cosine transform (DCT) coefficients, quantizing and encoding the coefficients, and using motion compensation techniques to encode most video frames as predictions from or between other frames. In particular, the encoded MPEG video stream is comprised of a series of groups of pictures (GOPs), and each GOP begins with an independently encoded (intra) I-frame and may include one or more following P-frames and B-frames. Each I-frame can be decoded

1 without information from any preceding and/or following frame. Decoding of a P-frame
 2 in general requires information from a preceding (I or P) frame in the same GOP.
 3 Decoding of a B-frame in general requires information both from a preceding (I or P)
 4 frame in the previous or the same GOP and a following (I or P) frame in the same GOP.

5 The MPEG-2 standard is documented in ISO/IEC International Standard (IS)
 6 13818-1, "Information Technology-Generic Coding of Moving Pictures and Associated
 7 Audio Information: Systems," ISO/IEC IS 13818-2, "Information Technology-Generic
 8 Coding of Moving Pictures and Associated Audio Information: Video," and ISO/IEC IS
 9 13818-3, "Information Technology-Generic Coding of Moving Pictures and Associated
 10 Audio Information: Audio," which are incorporated herein by reference. A concise
 11 introduction to MPEG is given in "A Guide to MPEG Fundamentals and Protocol
 12 Analysis (Including DVB and ATSC)," Tektronix Inc., 1997, incorporated herein by
 13 reference.

14 One application of MPEG-2 coded video is video-on-demand (VOD). In a VOD
 15 application, the video is stored in a server as MPEG-2 coded video. The server streams
 16 MPEG-2 coded video in real time to a subscriber's decoder. The subscriber may operate
 17 a remote control providing well-known classical videocassette recorder (VCR) functions
 18 including play, stop, fast-forward, fast-reverse, pause, slow-forward and slow-reverse.

19 Another application of MPEG-2 coded video is an MPEG-2 VCR. In an MPEG-2
 20 VCR application, the video is stored on a digital cassette in MPEG-2 coded video format.
 21 The MPEG-2 VCR streams MPEG-2 coded video in real time to an MPEG-2 decoder.
 22 The operator may operate a control providing well-known classical VCR functions
 23 including play, stop, fast-forward, fast-reverse, pause, slow-forward and slow-reverse.

1 A third application of MPEG-2 coded video is an MPEG-2 based video editing
2 station. In an MPEG-2 based video editing station, all video materials are stored in
3 MPEG-2 coded video format on tapes or disks. The operators may compile and edit the
4 MPEG-2 coded video in order to create a final broadcast version.

5 In the above applications of MPEG-2 coded video, it would be desirable to
6 provide an automatic method of detecting scene changes or identifying certain objects in
7 the (visual) scenes. For example, in lieu of a conventional fast-forward function, the
8 viewer could be provided with a function to skip forward to a next scene or skip back to a
9 previous scene, or a function to successively display new scenes in a forward or reverse
10 direction. Such scene display functions would omit the display of repetitious and
11 therefore irrelevant video frames in order for the viewer to find more quickly a new scene
12 from which regular speed forward-play may commence. The detection of a scene
13 change, however, involves a comparison of video information between successive
14 frames, and conventional methods for performing such a comparison are computationally
15 intensive. For real-time detection of scene changes, there is a need for fast,
16 computationally efficient, and reasonably successful detection of scene changes.

17

18 SUMMARY OF THE INVENTION

19 In accordance with a basic aspect of the invention, there is provided a method of
20 detecting edges in a compressed video sequence. The compressed video sequence
21 includes at least one frame of block encoded video data. The frame of block encoded
22 video data includes variable-length codes for transform coefficients of blocks of pixels in
23 the compressed video sequence. The transform coefficients include a respective DC

1 coefficient for each of the blocks of pixels. Each respective DC coefficient for at least
 2 some of the blocks of pixels is encoded as a respective variable-length code having a
 3 length indicating a certain range of differences in DC coefficient values between adjacent
 4 ones of the blocks of pixels. The method includes decoding only the length of the
 5 respective variable-length code for the respective DC coefficient for each of at least some
 6 of the blocks of pixels in order to produce an indication of whether or not the compressed
 7 video sequence includes an edge associated with each of the at least some of the blocks of
 8 pixels, and performing a code length threshold comparison upon the length of the
 9 respective variable-length code for the respective DC coefficient for each of the at least
 10 some of the blocks of pixels for producing at least one respective bit indicating whether
 11 or not the compressed video sequence includes an edge associated with said each of the at
 12 least some of the blocks of pixels.

13 In accordance with another aspect, the invention provides a method of detecting
 14 edges in a compressed video sequence. The compressed video sequence includes at least
 15 one I-frame of MPEG video data. The I-frame of MPEG video data includes variable-
 16 length codes for DCT coefficients of 8x8 pixel blocks in the image. The DCT
 17 coefficients include a respective DC coefficient for each of the 8x8 pixel blocks. Each
 18 respective DC coefficient for at least some of the 8x8 pixel blocks is encoded as a
 19 respective variable-length code having a length indicating a certain range of differences
 20 in DC coefficient values between adjacent ones of the 8x8 pixel blocks. The method
 21 includes decoding only the length of the respective variable-length code for the respective
 22 DC coefficient for each of the at least some of the 8x8 pixel blocks in order to produce an
 23 indication of whether or not the compressed video sequence includes an edge associated

1 with each of the at least some of the 8x8 pixel blocks, and performing a code length
 2 threshold comparison upon the length of the respective variable-length code for the
 3 respective DC coefficient for each of the at least some of the 8x8 pixel blocks for
 4 producing at least one respective bit indicating whether or not the compressed video
 5 sequence includes an edge associated with each of the at least some of the 8x8 pixel
 6 blocks.

7 In accordance with yet another aspect, the invention provides a method of
 8 detecting a scene change between I-frames of MPEG video data. The method includes
 9 detecting edges in images represented by the I-frames by decoding lengths of variable-
 10 length codes for DCT DC coefficients of 8x8 pixel blocks in the I-frames and performing
 11 code length threshold comparisons upon the decoded code lengths to produce respective
 12 edge indications for each of the I-frames. The method further includes comparing the
 13 edge indications between the I-frames in order to signal a scene change when there is a
 14 significant change in the edge indications between the I-frames.

15 In accordance with a final aspect, the invention provides a method of detecting a
 16 scene change between I-frames of MPEG video data. Each of the I-frames includes a
 17 series of 8x8 pixel blocks. The method includes detecting edges in images represented
 18 by the I-frames to produce a series of respective bits indicating whether or not an edge is
 19 associated with at least some of the 8x8 pixel blocks. The method further includes
 20 filtering the series of the respective bits indicating whether or not an edge is associated
 21 with the at least some of the 8x8 pixel blocks with a thinning filter in order to produce a
 22 filtered series of respective bits including more significant edge indications and excluding
 23 less significant edge indications. The method further includes operating a digital

processor to process the filtered series of respective bits in order to signal a scene change when there is a significant change in features between the I-frames.

BRIEF DESCRIPTION OF THE DRAWINGS

Other objects and advantages of the invention will become apparent upon reading the following detailed description with reference to the accompanying drawings, in which:

FIG. 1 is a block diagram of a system for processing of an MPEG coded video stream to provide a scene change signal upon reaching a current video I-frame having features that are substantially different from any features in the previous video I-frame;

FIG. 2 is a flow chart of a specific procedure that can be used for the processing of FIG. 1 to detect a scene change;

FIG. 3 is a diagram of logic for producing a luminance or chrominance edge indicating signal by comparing code lengths or values of the "dct_dc_size" luminance or chrominance attributes of each 8x8 pixel block to a threshold;

FIG. 4 is a diagram of logic for producing a combined luminance and chrominance edge indicating signal by comparing the sum of the code lengths or values of the "dct_dc_size" luminance attribute and "dct_dc_size" chrominance attributes of each 8x8 pixel block to a threshold length;

FIG. 5 is a matrix showing states of edge indicating bits produced from a frame of the MPEG-2 standard video test sequence named "Susie" by comparing the sum of the code lengths of the "dct_dc_size" luminance attribute and "dct_dc_size" chrominance

1 attributes of each 8x8 pixel block to a threshold length of 13 bits, and in which a bit
2 indicating an edge is shown as a black square in the matrix;

3 FIG. 6 is a logic diagram of a thinning filter responsive to a comparison of code
4 lengths between a current block and a prior block;

5 FIG. 7 is a first sheet of a flow chart of programming for a digital processor for
6 implementing the thinning filter of FIG. 6;

7 FIG. 8 is a second sheet of the flow chart begun in FIG. 7;

8 FIG. 9 is logic diagram of a thinning filter that adds logic to the thinning filter of
9 FIG. 6 in order to retain indications of less significant edges;

10 FIG. 10 is a first sheet of a flow chart of programming of a digital processor for
11 implementing the thinning filter of FIG. 9;

12 FIG. 11 is a second sheet of the flow chart begun in FIG. 10;

13 FIG. 12 is a logic diagram of circuitry for computing auto-coincidence counts and
14 cross-coincidence counts from frames of edge indicating data;

15 FIG. 13 is a flow chart showing the comparison of an auto-coincidence matrix for
16 edge indicating data of a current frame to an auto-coincidence matrix for edge indicating
17 data of a prior frame for detection of a scene change;

18 FIG. 14 is a logic diagram of circuitry for producing two separate channels of
19 edge indicating data, including a first channel of bits indicating edges having a negative
20 gradient in the horizontal direction, and a second channel of bits indicating edges having
21 a positive gradient in the horizontal direction;

22 FIG. 15 shows a block scanning sequence for (4:2:0) chrominance encoding;

23 FIG. 16 shows a block scanning sequence for (4:2:2) chrominance encoding;

1 FIG. 17 shows a block scanning sequence for (4:4:4, 4:2:2, and 4:2:0) luminance
2 and (4:4:4) chrominance encoding;

3 FIG. 18 is a diagram showing an edge with a positive or negative slope and a
4 positive horizontal gradient in pixel values between a current block and a prior block;

5 FIG. 19 is a diagram showing an edge with a positive or negative slope and a
6 negative horizontal gradient in pixel values between a current block and a prior block;

7 FIG. 20 is a diagram showing an edge with a positive or negative slope and a
8 positive horizontal gradient in pixel values within a prior block;

9 FIG. 21 is a diagram showing an edge with a positive or negative slope and a
10 negative horizontal gradient in pixel values within a prior block;

11 FIG. 22 is a diagram showing an edge with a positive or negative slope and a
12 positive horizontal gradient in pixel values within a current block;

13 FIG. 23 is a diagram showing an edge with a positive or negative slope and a
14 negative horizontal gradient in pixel values within a current block;

15 FIG. 24 shows logic responsive to attributes of only a current block for detecting
16 when it is more likely than not that there is an almost vertical edge in the current block
17 with a positive horizontal gradient component, and for detecting when it is more likely
18 than not that there is an almost vertical edge in the current block with a negative
19 horizontal gradient component;

20 FIG. 25 shows logic responsive to attributes of only a current block for detecting
21 when it is more likely than not that there is an almost horizontal edge in the current block
22 with a positive vertical gradient component; and for detecting when it is more likely than

1 not that there is an almost horizontal edge in the current block with a negative vertical
2 gradient component; and

3 FIG. 26 is a flow chart of a procedure for estimating the gradient vector of an
4 edge.

5 While the invention is susceptible to various modifications and alternative forms,
6 specific embodiments thereof have been shown by way of example in the drawings and
7 will be described in detail. It should be understood, however, that it is not intended to
8 limit the form of the invention to the particular forms shown, but on the contrary, the
9 intention is to cover all modifications, equivalents, and alternatives falling within the
10 scope of the invention as defined by the appended claims.

11

12 DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS

13 With reference to FIG. 1, there is shown a block diagram of a system for
14 processing of an MPEG coded video stream to provide a scene change signal upon
15 reaching a current video frame having features that are substantially different from
16 features in the previous video frame. The features are comprised of edges. An edge
17 detector 20 detects the edges occurring between and within blocks in each frame. A
18 thinning filter 21 receives edge indications from the edge detector 20, and eliminates
19 some of the less significant edge indications.

20 The thinning filter in effect refines the initial edge information and produces
21 information that can be visualized as edge graph information. The amount of the
22 thinning should be carefully selected, and a three-pixel deep thinning is preferred. The
23 remaining edge indications are loaded into a selected one of two ping-pong (parallel)

1 frame buffers 22 and 23.

2 In order to provide a scene change signal, a digital processor 24 accesses edge
3 indications for a current frame in a selected one of the frame buffers and edge indications
4 for a prior frame in the other one of the frame buffers. The digital processor 24 performs
5 feature extraction and comparison and detects a scene change when there is a significant
6 difference between features in the current frame and features in the previous frame.
7 When a scene change is detected, the current frame is identified as a "key frame" in the
8 MPEG video sequence. If the MPEG video sequence is being read from or written to an
9 MPEG file in data storage, the frames identified as "key frames" could be marked as such
10 in the data storage in a frame or GOP index table or in a metadata file associated with the
11 MPEG file.

12 There are many known edge detection procedures having varying complexities.
13 However, most are quite complex and computationally expensive. For real-time
14 applications, the computation time is a major factor. In addition not all applications need
15 edge detection results with high fidelity. For the edge detector 20 in the system of FIG.
16 1, there is a need for a fast and computationally efficient but reasonably successful edge
17 detection procedure.

18 As will be further described below with reference to FIGS. 3 and 4, the preferred
19 edge detection procedure for the MPEG scene change detection system of FIG. 1 is based
20 on the DC coefficients rather than all 64 coefficients within the DCT block. Therefore,
21 the resolution of the edge detail image drops to $1/8^{\text{th}}$ in both horizontal and vertical
22 directions for a resultant reduction to $1/64$ times in the resolution. Moreover, the
23 preferred procedure also does not need to decode the actual value of the DC coefficient.

1 Instead, the edge detection procedure is responsive to the length of the variable-length
2 code used for the predictive encoding of the DC coefficient (dct_dc_differential).

3 FIG. 2 is a flow chart of a specific procedure that can be used by the system of
4 FIG. 1 to detect a scene change in the MPEG video sequence. In a first step 31, the
5 lengths of the "dct_dc_size" MPEG variable-length codes or alternatively the values of
6 the "dct_dc_size" attributes are compared to at least one length threshold to detect edges
7 in the video frame. The values of the "dct_dc_size" attributes may be employed when the
8 lengths of the "dct_dc_size" variable length codes do not uniquely identify the magnitude
9 class of the differential DC coefficient value. In MPEG-2, for example, the values of the
10 DCT DC coefficients for the first 8x8 luminance (Y) and chrominance (Cb and Cr) pixel
11 blocks at the beginning of each slice of macroblocks in the frame are coded relative to a
12 fixed value, whereas the DCT DC coefficient for each following 8x8 pixel block in the
13 slice is coded relative to the value of the previous DC coefficient in the same slice and in
14 the same color component. In particular, the DCT DC coefficient for each following 8x8
15 pixel block in the slice is coded as a difference relative to the previous DCT DC
16 coefficient in the same color component. The relative ordering of the 8x8 pixel blocks is
17 defined on the basis of a block scan order described in the MPEG-2 standard. A
18 substantial difference between the DCT DC coefficient values of a pair of neighboring
19 blocks in the same color component, implies that the blocks have different color
20 (luminance and/or chrominance) features, and hence an edge is likely to be there.

21 The following Tables B-12 and B-13, reproduced from page 125 of the above-
22 cited MPEG-2 standard document ISO/IEC IS 13818-2, show how the respective
23 differences between the DCT DC coefficients of neighboring 8x8 pixel blocks of the

1 same color component, luminance (Y) or chrominance (Cb or Cr), are variable-length
 2 coded. The tables are such that as the difference becomes larger in magnitude, more bits
 3 are required to code it. Therefore, the longer VLC codes indicate substantial differences,
 4 and the VLC code length can be compared to a threshold length to provide an indication
 5 of whether or not an edge with a contrast (gradient) stronger than a predefined value is
 6 present in the vicinity of the respective 8x8 pixel bock in the image.

7

8 Table B-12 --- Variable length codes for dct_dc_size_luminance

9	Variable length code	dct_dc_size_luminance
10	100	0
11	00	1
12	01	2
13	101	3
14	110	4
15	1110	5
16	1111 0	6
17	1111 10	7
18	1111 110	8
19	1111 1110	9
20	1111 1111 0	10
21	1111 1111 1	11

22

23 Table B-13 --- Variable length codes for dct_dc_size_chrominance

1	Variable length code	dct_dc_size_chrominance
2	00	0
3	01	1
4	10	2
5	110	3
6	1110	4
7	1111 0	5
8	1111 10	6
9	1111 110	7
10	1111 1110	8
11	1111 1111 0	9
12	1111 1111 10	10
13	1111 1111 11	11
14		

15 To detect edges, the MPEG-2 video bit stream is only parsed and the dct_dc_size
16 (luminance and/or chrominance) variable length codes possibly simply inspected, to
17 determine how many bits are used to code the difference between the DCT DC
18 coefficient of the current block and the DCT DC coefficient of the previous block of the
19 same color component (without a need for finding out what the current DC coefficient
20 value is nor what the difference is). This code length is compared to a threshold length to
21 produce a bit indicating the presence or absence of an edge, Case I: code length greater
22 than or equal to the threshold, and Case II: code length less than the threshold,
23 respectively. There is a trade-off in selecting the threshold length. If the threshold length

1 is too large, only the strongest edges will be detected. If the threshold length is too small,
 2 some features will mistakenly be detected as edges, i.e., the false alarm rate will increase,
 3 especially in high frequency texture areas (for example, the hair in the standard MPEG
 4 video test sequence known as "Susie"). A mid-range of values about 7, 8, or 9 for the
 5 luminance or chrominance dct_dc_sizes is a good compromise.

6 The edge information obtained from the threshold comparison of step 31 may not
 7 be immediately suitable for many applications, because a very large change in the value
 8 of the DC coefficient may occur over a sequence of consecutive blocks. Such a change
 9 appears in the image as a single edge feature yet the threshold comparison may produce a
 10 sequence of consecutive edge indications. In such a case it is desired to confine the edge
 11 indication to the location of the maximum gradient, i.e., to the location of the maximum
 12 rate of color (luminance and/or chrominance) change, which is at about the beginning of
 13 the block having the largest dct_dc_size value, or the longest code length. Therefore, in
 14 step 32, the edge indicating signal is applied to a thinning filter. The thinning filter
 15 compares the lengths of the dct_dc_size variable-length codes or alternatively the values
 16 of the dct_dc_size attributes for adjacent blocks. Such a thinning filter is further
 17 described below with reference to FIGs. 6 to 8. In the case of two edges that are
 18 separated by about the width or height of one 8x8 pixel block, such a thinning filter will
 19 also select the most significant of the two edges. The elimination of the less significant
 20 edge indications produces a frame of bits that is more representative of the distinctive
 21 features in the image.

22 The frame of bits output from the thinning filter can be further processed to
 23 produce auto-coincidence counts that are characteristics of the individual frames, and

1 cross-coincidence counts that indicate the similarity between the contents of the current
 2 frame and the contents of other frames. The auto-coincidence counts and cross-
 3 coincidence counts are computed in a fashion similar to auto-correlation and cross-
 4 correlation. Correlation involves a summing of products, or in the case of binary signals,
 5 a summing of logic 1's produced by an exclusive-NOR function. In contrast, coincidence
 6 involves a summing of logic 1's produced by a logic AND function. In other words,
 7 what is significant for coincidence is a match of logic 1's.

8 One way of detecting a scene change is to compute a coincidence coefficient
 9 indicating a degree of coincidence between significant edges in a current frame and
 10 significant edges in a prior frame to within a distance of a small number of (8x8 pixel)
 11 blocks, such as 0 or 1 block. In general the number of blocks can be determined on the
 12 basis of the temporal distance between the current and the prior frames and the amount of
 13 motion in the scene in that temporal vicinity. For example, in step 33, an auto-
 14 coincidence count (C_{cc}) is computed for the current frame that indicates the number of
 15 significant edges that match between the current frame and a copy of it without any shift
 16 and between the current frame and its copies shifted by one block position in each of the
 17 two horizontal directions, the two vertical directions, and the four diagonal directions. A
 18 way of computing such an auto-coincidence count is described below with reference to
 19 FIG. 12. In step 34, a cross-coincidence count (C_{cp}) is computed that indicates the
 20 number of significant edges that match between the current frame and the prior frame
 21 without any shift and between the current frame and the copies of the prior frame shifted
 22 by one block position in each of the two horizontal directions, the two vertical directions,
 23 and the four diagonal directions. A way of computing such a cross-coincidence count is

1 described below with reference to FIG. 12.

2 In step 35, the coincidence coefficient is computed by dividing the cross-
3 coincidence count (C_{cp}) by the arithmetic mean of the auto-coincidence count (C_{cc}) for
4 the current frame and the auto-coincidence count (C_{pp}) for the prior frame. In step 36, the
5 coincidence coefficient is compared to at least one coincidence threshold to detect a
6 scene change when the coincidence coefficient fails to reach the threshold. For example,
7 the coincidence coefficient ranges from zero, indicating no coincidence between
8 significant edges in the current and prior frames, to one, indicating complete coincidence
9 between significant edges in the current and prior frames, and the coincidence threshold
10 is about 0.3 to 0.7.

11 With reference to FIG. 3, there is shown logic for producing a luminance or
12 chrominance edge signal. If the current block is the first block in the slice, then an
13 inverter 41 and an AND gate 42 set the edge signal to a logic zero. Otherwise, a decoder
14 43 decodes the code length of the `dct_dc_size_luminance` or `dct_dc_size_chrominance`
15 variable-length code. A comparator 44 compares the code length to a length threshold to
16 produce the luminance or chrominance edge signal. The comparator 44 produces a logic
17 1 when the code length is equal to or greater than the length threshold.

18 For a color image, the logic of FIG. 3 produces three relatively independent edge
19 signals, one for luminance and one for each of the two chrominance components Cb and
20 Cr. In MPEG-2, the chrominance components of the video frames typically are sub-
21 sampled with respect to luminance, so that there will be fewer chrominance edge
22 indicating bits than luminance edge indicating bits. However, as shown in FIG. 4, it is
23 possible to produce a combined luminance and chrominance edge signal in order to

1 reduce the number of edge indicating bits. The logic of FIG. 4, for example, produces a
2 number of combined luminance and chrominance edge indicating bits equal to the
3 number of luminance edge indicating bits produced by the logic of FIG. 3. For this
4 purpose, depending on the chrominance sub-sampling pattern, each
5 dct_dc_size_chrominance code length from both Cb and Cr components may have to be
6 repeatedly used with multiple dct_dc_size_luminance code lengths.

7 As shown in FIG. 4, an inverter 51 and an AND gate 52 set the luminance and
8 chrominance edge signal to a logic zero for the first block in every slice. Otherwise,
9 decoders 53, 54 and 55 decode the code length of the dct_dc_size_luminance variable-
10 length code and the code lengths of the dct_dc_size_chrominance variable-length codes
11 for both Cb and Cr components. An adder 56 adds the luminance code length to the
12 chrominance code lengths to compute a combined code length. A comparator 57
13 compares the combined code length to a threshold length to produce the combined
14 luminance and chrominance edge signal. The comparator 57 produces a logic 1 when the
15 combined code length is equal to or greater than the length threshold.

16 Although the logic in FIG. 4 combines the luminance and chrominance code
17 lengths by a simple addition 56, it should be apparent that the luminance and
18 chrominance code lengths could be combined in various other ways to produce a
19 combined code length, such as by computing a weighted arithmetic mean, a geometric
20 mean, a sum of the squares of the luminance and chrominance code lengths, or some
21 other average or combination of the luminance and/or chrominance code lengths.

22 FIG. 5 shows a frame of the edge signal produced by the logic of FIG. 4 on a
23 frame from the MPEG-2 standard video test sequence named "Susie" and a threshold

1 length of 13 bits. A black square in a cell of the matrix in FIG. 5 indicates that the edge
2 signal is a logic one for the 8x8 pixel block associated with that cell.

3 FIG. 6 is a logic diagram of a thinning filter responsive to a comparison of the
4 dct_dc_size code lengths between a current block and a prior block. The code length
5 from the prior block is held in a register 71, and a comparator 72 provides a first binary
6 output 73 that is a logic one only when the code length from the current block is less than
7 the code length from the prior block, and a second binary output 74 that is a logic one
8 only when the code length from the current block is greater than the code length from the
9 prior block. The thinning filter is also responsive to the edge signal for the current block,
10 and a single-bit register 75 holds the edge signal for the prior block. The thinning filter
11 provides a filtered edge signal 76 from the output of an AND gate 77.

12 The filtered edge signal 76 is a logic one when an edge is indicated for the prior
13 block and the prior block has a dct_dc_size code length that is not less than the
14 dct_dc_size code length for the current block if an edge is indicated for the current block,
15 and the prior block has a dct_dc_size code length that is not less than the dct_dc_size
16 code length for the block just before the prior block if an edge is indicated for the block
17 just before the prior block. An inverter 78 and an OR gate 79 provide a signal to the
18 AND gate 77 indicating whether or not the prior block has a dct_dc_size code length that
19 is not less than the dct_dc_size code length for the current block if an edge is indicated
20 for the current block. An inverter 80, an OR gate 81, and a single-bit register 82 provide
21 a signal to the AND gate 77 indicating whether or not the prior block has a dct_dc_size
22 code length that is not less than the dct_dc_size code length for the block just before the
23 prior block if an edge is indicated for the block just before the prior block.

1 FIG. 7 is a first sheet of a flow chart of a program for a digital processor to
 2 implement the thinning filter of FIG. 6. In a first step 91, the processor branches to step
 3 92 if a logical variable (i.e. a Boolean variable) "PRIOR_EDGE" (representing the output
 4 of the single-bit register 75 in FIG. 6) is zero. In step 92, the processor sets the output
 5 signal "FILTERED_EDGE" (representing the output of the AND gate 77 in FIG. 6) to
 6 zero. In step 93, the processor sets a logical variable "DAGTB" (representing the input
 7 to the single-bit register 82 in FIG. 6) to one, and execution continues in step 105 in FIG.
 8 8. In step 105, the value of a logical variable "QAGTB" (representing the output of the
 9 single-bit register 82 in FIG. 6) is set equal to the value of the logical variable "DAGTB"
 10 and execution continues to step 94. In step 94, a variable "PRIOR_CDLENG"
 11 (representing the output of the register 71 in FIG. 6) is set to the value of "CDLENG"
 12 (representing the input to the register 71 in FIG. 6). In step 95, the logical variable
 13 "PRIOR_EDGE" is set to the value of a variable "EDGE" (representing the edge signal
 14 in FIG. 6, as produced by the edge detector of FIG. 3 or the edge detector of FIG. 4 and
 15 connected to the input of the single-bit register 75 in FIG. 6). After step 95, the
 16 computations of the thinning filter are done for the processing of the "EDGE" signal and
 17 the "CDLENG" signal for the current block.

18 In FIG. 7, step 91, if the processor finds that the "PRIOR_EDGE" variable is not
 19 equal to zero, execution continues to step 96. In step 96, the processor compares the
 20 value of the variable "CDLENG" (representing the dct_dc_size code length for the
 21 current block) to the value of the variable "PRIOR_CDLENG" (representing the
 22 dct_dc_size code length for the prior block). If the comparison finds that "CDLENG" is
 23 greater than "PRIOR_CDLENG," then execution branches from step 97 to step 98, to set

1 the logical variable "DAGTB" to one. Otherwise, execution continues from step 97 to
2 step 99, to set the logical variable "DAGTB" to zero. After steps 98 and 99, execution
3 continues to step 100.

4 In step 100, execution continues to step 101 if the comparison finds that
5 "CDLENG" is not less than "PRIOR_CDLENG." In step 101, if the edge signal
6 "EDGE" is equal to zero, then execution branches to step 102. Execution also branches
7 to step 102, from step 100, if the comparison finds that "CDLENG" is less than
8 "PRIOR_CDLENG." In step 102, execution branches to step 103 if the logical variable
9 "QAGTB" is equal to zero. Execution also continues from step 101 to step 103 if the
10 edge signal "EDGE" is not equal to zero. In step 103, the output signal
11 "FILTERED_EDGE" is set to zero, and execution continues to step 105 in FIG. 8. In
12 step 102, if the logical variable "QAGTB" is not equal to zero, the execution branches to
13 step 104. In step 104, the output signal "FILTERED_EDGE" is set to one, and execution
14 continues to step 105 in FIG. 8. In step 105, the value of the logical variable "QAGTB"
15 is set equal to the value of the logical variable "DAGTB," and execution continues to step
16 94. In step 94, the variable "PRIOR_CDLENG" is set to the value of the variable
17 "CDLENG" and execution continues to step 95. In step 95, the logical variable
18 "PRIOR_EDGE" is set to the value of the variable "EDGE". After step 95, the
19 computations of the thinning filter are done for the processing of the "EDGE" and
20 "CDLENG" signals for the current block along this alternate path of execution taken
21 when "PRIOR_EDGE" does not compare equal to 0 at step 91.

22 FIG. 9 is a logic diagram of a thinning filter that adds logic to the thinning filter of
23 FIG. 6 in order to retain indications of less significant edges. Components 101 to 112 in

FIG. 9 further includes a single-bit register 113 that receives the sign of either the DCT DC luminance or chrominance (Cb or Cr) for the current block, and provides the sign of the DCT DC luminance or chrominance (Cb or Cr) of the prior block to an exclusive OR gate 114. If the edge signal received by the single-bit register 105 indicates luminance edges, then the single bit register 113 should receive the sign of the DCT DC luminance. If the edge signal received by the single-bit register 105 indicates chrominance (Cb or Cr) edges, then the single-bit register 113 should receive the sign of the DCT DC chrominance (Cb or Cr respectively). If the edge signal received by the single-bit register 105 indicates combined luminance and chrominance edges (such as the output of the edge detector in FIG. 4), then the single-bit register 113 should receive either the sign of the DCT DC luminance, or the sign of a DCT DC chrominance (Cb or Cr). When the register 113 is receiving the sign of the DCT DC luminance, the exclusive-OR gate 114 produces a logic 1 when there has been a change in the sign of the DCT DC luminance. When the register 113 is receiving the sign of a DCT DC chrominance (Cb or Cr) , the exclusive-OR gate 114 produces a logic 1 when there has been a change in the sign of the same DCT DC chrominance component. In either case, when there has been such a change in sign, the logic 1 output from the exclusive-OR gate 114 is effective to inhibit any comparison from the comparator 102 from rejecting an edge indication during the filtering process. In particular, the logic 1 output from the exclusive-OR gate 114 drives the outputs of the OR gate 109 and the OR gate 111 both to logic 1.

FIG. 10 is a first sheet of a flow chart of a program for a digital processor to implement the thinning filter of FIG. 9. Steps 121 to 123 are similar to steps 91 to 93 in

FIG. 7. Execution continues from step 123 to step 140 in FIG. 11. Steps 140 and 124 to 125 in FIG. 11 are similar to steps 105 and 94 to 95 in FIG. 8. Step 126 in FIG. 11 sets a logical variable "PRIOR_SIGN" (representing the output of the single-bit register 113 in FIG. 9) to the value of a logical variable "SIGN" (representing the input to the single-bit register 113 in FIG. 9) indicating the sign of the dct_dc_luminance or dct_dc_chrominance (for Cb or Cr components), in order to implement the single-bit register 113 of FIG. 9. After step 126, execution returns.

In step 121 of FIG. 10, if the value of the logical variable "PRIOR_EDGE" is not equal to zero, then execution continues to step 127. In step 127, the processor computes the exclusive-OR of the sign of the dct_dc_luminance or dct_dc_chrominance (for Cb or Cr components) and the value of the logical variable "PRIOR_SIGN," in order to implement the exclusive-OR gate 114 in FIG. 9. In step 128, if the exclusive-OR result is not equal to zero, execution branches to step 129 to set the logical variable "DAGTB" equal to 1. After step 129, execution continues to step 137 in FIG. 11. Steps 137 to 140 in FIG. 11 are similar to steps 102 to 105 in FIGs. 7 and 8.

In step 128, if the exclusive-OR result is equal to zero, then execution continues to step 130. Steps 130, 132, 133 and 134 are similar to steps 96 to 99 of FIG. 7. After steps 133 or 134, execution continues to step 135 in FIG. 11. Steps 135 to 140 are similar to steps 100 to 105 in FIGs. 7 and 8.

FIG. 12 is a logic diagram of a circuitry for computing auto-coincidence counts and cross-coincidence counts from frames of edge indicating data. The coincident count is computed between a frame of data in a first frame buffer 151 and a frame of data in a second frame buffer 152. For a cross-coincident count, the data in the first frame buffer

1 151 are different from the data in the second frame buffer 152. For an auto-coincidence
2 count, the data in the first frame buffer 151 are the same as the data in the second frame
3 buffer 152.

4 The circuitry of FIG. 12 successively compares selected bytes or words of data
5 from the first frame buffer 151 to selected bytes or words of data from the second frame
6 buffer 152. A first multiplexer 153 selects an addressed byte or word from the first frame
7 buffer 151 and loads the addressed byte or word into a first (cyclic) shift register 154. A
8 second multiplexer 155 selects an addressed byte or word from the second frame buffer
9 152 and loads the addressed byte or word into a second (cyclic) shift register 156. An
10 array of AND gates 157 compare bits of a byte or word in the first shift register 154 to
11 corresponding bits of a byte or word in the second shift register 156. A population
12 counter 158 computes the number of logic 1's produced by the AND gates. The
13 population counter 158, for example, is comprised of an array of full adders, or a memory
14 programmed with a look-up table. An adder 159 adds the number from the population
15 counter to the output of accumulator register 160 to accumulate a coincidence count.

16 The circuitry of FIG. 12 can be implemented by suitable programming of a
17 general purpose digital processor. If the general purpose digital processor does not have
18 a machine instruction providing a population count function, then the population count
19 function can be implemented by a look-up table in random access memory. A 256 entry
20 look-up table with 4 bits of storage per entry, is sufficient for computing the population
21 count of a byte of data from the logical AND of a single byte in the first shift register 154
22 to a single byte in the second shift register 156. Following is a pseudo-code listing for
23 the programming of a general-purpose digital processor for computing the auto-

1 coincidence counts of step 33 in FIG. 2 or the cross-coincidence count of step 34 in FIG.
 2 2. For the auto-coincidence count of step 33, FRAME1 is the same as FRAME2. For the
 3 cross-coincidence count of step 34, FRAME1 is different from FRAME2.

4
 5 /* FRAME HAS "M" ROWS AND "N" COLUMNS OF 8X8 BLOCKS */
 6 /* THE EDGE INDICATING BIT FOR THE FIRST BLOCK IN EACH */
 7 /* ROW IS LEFT-JUSTIFIED TO THE LEAST SIGNIFICANT BIT */
 8 /* POSITION. THE FRAME BUFFER BYTE CONTAINING THE */
 9 /* EDGE INDICATING BIT FOR THE LAST BLOCK IN EACH ROW */
 10 /* HAS ITS MOST SIGNIFICANT BITS ZERO-FILLED IF THE */
 11 /* NUMBER OF BLOCKS PER ROW IS NOT A MULTIPLE OF */
 12 /* EIGHT. THEREFORE THE FRAME BUFFER HAS "M" ROWS */
 13 /* OF BYTES AND "P" COLUMNS OF BYTES WHERE "P" IS */
 14 /* THE NUMBER OF BYTES PER ROW COMPUTED AS: */
 15 /* */
 16 /* P = INT(N/8) */
 17 /* IF (REM(N/8).NE.0) THEN P = P + 1 */
 18 /* WHERE THE INT(.) FUNCTION RETURNS THE INTEGER PART */
 19 /* OF ITS FLOATING POINT ARGUMENT AND THE FUNCTION */
 20 /* REM(.) RETURNS THE REMAINDER OF ITS ARGUMENT */
 21 /* WHICH IS IN THE FORM OF A RATIO OF INTEGERS. */
 22
 23 /* BEGIN COINCIDENCE COUNT CALCULATION */

```

1
2  CLEAR ACCUMULATOR (OF SIZE ( INT( LOG_2(9*M*P) ) + 1) BITS)
3  FOR I = 2 TO M-1
4    FOR K = 1 TO 3
5      CLEAR CYCLIC SHIFT_REGISTER_1 (OF SIZE 9 BITS)
6      CLEAR CYCLIC SHIFT_REGISTER_2 (OF SIZE 9 BITS)
7      FOR J = 1 TO P
8
9        BYTE_ADDR_1 = ( ( I - 1 ) * P ) + J
10       BYTE_ADDR_2 = ( ( I + K - 3 ) * P ) + J
11
12      /* LOAD SHIFT_REGISTER_1 AND SHIFT_REGISTER_2 */
13      /* IN THE FOLLOWING LOAD OPERATIONS ALL 9 BITS OF THE */
14      /* SHIFT REGISTER UNDER CONSIDERATION ARE LOADED; */
15      /* THE LEFTMOST (LEAST SIGNIFICANT) 8 BITS ARE LOADED */
16      /* WITH THE ADDRESSED BYTE; AND THE REMAINING */
17      /* RIGHTMOST (MOST SIGNIFICANT) 1 BIT IS LOADED WITH */
18      /* THE LEAST SIGNIFICANT BIT OF THE NEXT BYTE ON THE */
19      /* SAME ROW */
20
21     SHIFT_REGISTER_1 = FRAME1(BYTE_ADDR_1)
22     SHIFT_REGISTER_2 = FRAME2(BYTE_ADDR_2)
23

```

```
1  /* DO HORIZONTALLY ALIGNED COMPARISON */
2  /* IN ALL CONSEQUENT LOGICAL "AND" OPERATIONS, ONLY */
3  /* THE LEFTMOST (LEAST SIGNIFICANT) 8 BITS OF BOTH */
4  /* REGISTERS ARE INVOLVED TO PRODUCE AN 8 BIT RESULT */
5  /* (INDEX) */
6
7  INDEX = SHIFT_REGISTER_1.AND.SHIFT_REGISTER_2
8  INCREMENT = POPULATION_COUNT_FUNCTION(INDEX)
9  ACCUMULATOR = ACCUMULATOR + INCREMENT
10
11 /* DO COMPARISON WITH ONE BLOCK HORIZONTAL SHIFT OF */
12 /* FRAME 1 TO THE LEFT WITH RESPECT TO FRAME 2 */
13
14 CYCLICLY ROTATE SHIFT_REGISTER_1 LEFT BY ONE BIT
15 INDEX = SHIFT_REGISTER_1.AND.SHIFT_REGISTER_2
16 INCREMENT = POPULATION_COUNT_FUNCTION(INDEX)
17 ACCUMULATOR = ACCUMULATOR + INCREMENT
18
19 /* RESTORE SHIFT_REGISTER_1 TO ORIGINAL STATE */
20
21 CYCLICLY ROTATE SHIFT_REGISTER_1 RIGHT BY ONE BIT
22
23 /* DO COMPARISON WITH ONE BLOCK HORIZONTAL SHIFT OF */
```

```

1  /* FRAME2 TO THE LEFT WITH RESPECT TO FRAME 1          */
2
3  CYCLICLY ROTATE SHIFT_REGISTER_2 LEFT BY ONE BIT
4  INDEX = SHIFT_REGISTER_1.AND.SHIFT_REGISTER_2
5  INCREMENT = POPULATION_COUNT_FUNCTION(INDEX)
6  ACCUMULATOR = ACCUMULATOR + INCREMENT
7
8  /* RESTORE SHIFT_REGISTER_2 TO ORIGINAL STATE */
9
10 CYCLICLY ROTATE SHIFT_REGISTER_2 RIGHT BY ONE BIT
11
12 NEXT J
13 NEXT K
14 NEXT I

```

16 If it is desired to detect changes between the contents of frames regardless of
 17 possibly extensive (more than one 8x8 pixel block in all possible horizontal, vertical and
 18 diagonal directions) translations of objects in the frames, it may be more desirable to
 19 compute the variance between the auto-coincidence matrices of the frames rather than the
 20 cross-coincidence counts between the frames. The auto-coincidence matrix of a frame is
 21 a matrix of auto-coincidence counts for the frame. The auto-coincidence count in the cell
 22 (m, n) of the auto-coincidence matrix is the count of matches between the frame and the
 23 same frame shifted by m rows and n columns. The counts in the auto-coincidence matrix

1 for limited ranges of m and n , for example, $m = 0, 1, 2, \dots, M$ and $n = 0, 1, 2, \dots, N$ for
 2 relatively small values of M and N such as $M = 3$ and $N = 3$ provide a way of identifying
 3 the frame. These counts are relatively independent of the translation of the frame as a
 4 whole and are also relatively independent of the translation of objects in the frame so
 5 long as the objects do not completely overlap each other. The counts could be computed
 6 for each I-frame of an MPEG-2 clip and stored as part of the metadata of the clip.

7 FIG. 13 is a flow chart showing the comparison of an auto-coincidence matrix for
 8 edge indicating data of a current frame to an auto-coincidence matrix for edge indicating
 9 data of a prior frame for detection of a scene change. In a first step 171, the processor
 10 computes the auto-coincidence matrix $C_{c(M,N)}$ for the current frame. In step 172, the
 11 processor computes a coefficient of variance between the auto-coincidence matrix $C_{c(M,N)}$
 12 for the current frame and the auto-coincidence matrix $C_{p(M,N)}$ for the prior frame. In step
 13 173, the processor saves the auto-coincidence matrix of the current frame as the auto-
 14 coincidence matrix for the prior frame in order to be used when processing the next
 15 frame. In step 174, the processor compares the coefficient of variance between the auto-
 16 coincidence matrices to a threshold to detect a scene change when the coefficient of
 17 variance is equal to or greater than the threshold. For example, the coefficient of
 18 variance varies from 0 for no variance to 1 for complete variance, and the threshold is
 19 about 0.2 to 0.6.

20 Unless there is a scarcity of computational resources, it is desirable to use
 21 information about the gradient or orientation of the edges in order to better distinguish a
 22 scene change from change in frames due to some motion of the background or objects in
 23 a frame or due to changes in lighting conditions between frames. It is most desirable to

1 provide additional independent bits or channels of bits in such a way that the frame of the
2 bits for each channel will be relatively sparse. Under these conditions, there will be a
3 greater variance between the auto-coincidence matrix elements of different frames for
4 relatively small values of m and n.

5 One way of distinguishing edges is based on the signs of the horizontal and/or
6 vertical components of the gradient (vector) for each edge, as indicated by the sign of the
7 differential DCT DC (C_{00}) coefficient and the block scan direction reaching the block
8 containing this coefficient. For example, for a luminance edge, it is the sign of the DCT
9 C_{00} coefficient for luminance and the direction along which the prediction is performed to
10 produce this coefficient; for a chrominance edge, it is the sign of the DCT C_{00} coefficient
11 for the same chrominance channel and the prediction direction for this coefficient; and
12 for a combination luminance and chrominance edge, it is the sign of the DCT C_{00}
13 coefficient for luminance and the prediction direction for this coefficient. In any of these
14 cases, the sign is very distinctive because the magnitude of the DCT C_{00} coefficient, as
15 classified by the value of "dct_dc_size," is at least as large as the threshold for edge
16 detection.

17 FIG. 14 shows how the sign of the DCT C_{00} coefficient and the direction in which
18 it was predicted can be used to split the thinning filtered edge signal into two independent
19 channels, one providing bits indicating edges with a negative horizontal gradient
20 component, and one providing bits indicating edges with a positive horizontal gradient
21 component. An exclusive-OR gate 180 produces a logical signal representing the sign of
22 the horizontal component of the local edge gradient ("0" for positive and "1" for
23 negative) based on its two inputs SIGN_DCT_ C_{00} and PD_DCT_ C_{00} , respectively the

1 sign and the prediction direction of the current block's differential DCT DC coefficient.
 2 The logical value of PD_DCT_ C₀₀ depends on block position in one of three scan
 3 sequences for the particular chrominance sub-sampling method, as described further
 4 below with reference to FIGS. 15 to 17. A single-bit register 181 provides a delay of the
 5 sign of the horizontal gradient component corresponding to the delay of the edge signal
 6 through the thinning filter 182, which is the thinning filter of FIG. 6 or the thinning filter
 7 of FIG. 9. An AND gate 183 combines the delayed sign of the horizontal gradient
 8 component from the single-bit register 181 with the thinning filtered edge signal from the
 9 thinning filter 182 to produce a signal indicating edges with a negative horizontal
 10 gradient component. An inverter 184 inverts the delayed sign of the horizontal gradient
 11 component from the single-bit register 181. An AND gate 185 combines the output of
 12 the inverter 184 with the thinning filtered edge signal from the thinning filter 182 to
 13 produce a signal indicating edges with a positive horizontal gradient component.

14 FIG. 15 shows how the logical value of PD_DCT_ C₀₀ depends on block position
 15 in the scan sequence for (4:2:0) chrominance (Cb and Cr). In this case, the scan order is a
 16 simple raster scan, and the value of PD_DCT_ C₀₀ is a constant logic zero.

17 FIG. 16 shows how the logical value of PD_DCT_ C₀₀ depends on block position
 18 in the scan sequence for (4:2:2) chrominance (Cb and Cr). In this case, the scan order is a
 19 simple saw-tooth pattern. The value of PD_DCT_ C₀₀ is a sequence of 101010... .

20 FIG. 17 shows how the logical value of PD_DCT_ C₀₀ depends on block position
 21 in the scan sequence for (4:4:4, 4:2:2, and 4:2:0) luminance and (4:4:4) chrominance (Cb
 22 and Cr). In this case, the scan order is rather complex, including one step backward
 23 (where PD_DCT_ C₀₀ has a value of logic 1) for every three steps forward (where

1 PD_DCT_ C_{00} has a value of logic 0). In other words, the value of PD_DCT_ C_{00} is a
 2 sequence of 0100010001000100....

3 Another source of additional information about the edges is the sign and
 4 magnitude of the lowest (horizontal and vertical) frequency DCT AC coefficients.
 5 Vertically and horizontally oriented edges depending on their contrast levels, respectively
 6 induce horizontal and vertical frequency components. For example, FIG. 18 shows an
 7 idealized case of a close-to-vertical edge with a positive horizontal gradient component
 8 (pixel values increasing towards right) being detected by an edge signal that is a logic 1
 9 for a current block, and FIG. 19 shows an idealized case of another close-to-vertical edge
 10 with a negative horizontal gradient component (pixel values decreasing towards right)
 11 being detected by an edge signal that is a logic 1 for a current block. In each of these
 12 idealized (due to their involving an ideal step-edge) cases, the edge is centered almost
 13 right at the vertical boundary between the current block and the prior block. The
 14 discontinuity of the edge is predominantly confined to neither the current block nor the
 15 prior block. Therefore the contents of both the current block and the prior block are
 16 predominantly low-pass (smooth) and the existence of a vertical edge between the blocks
 17 for the case illustrated in FIG. 18 and FIG. 19 will be revealed only through a significant
 18 change in the mean pixel value from the prior to the current block or equivalently a large
 19 magnitude differential DCT DC (C_{00}) coefficient in the current block. As long as the
 20 edges of FIG. 18 and FIG. 19 are close to being vertical and located almost at the vertical
 21 boundary between the prior and current blocks, the exact values of their slopes are
 22 immaterial to the above observation. As the almost vertical edges of FIG. 18 or FIG. 19
 23 shift into either the prior block or the current block as illustrated in FIGS. 20 and 21 or

1 FIGS. 22 and 23, respectively, the difference in the mean pixel values of the prior and the
 2 current block or equivalently the magnitude of the differential DCT DC (C_{00}) coefficient
 3 in the current block decreases. To detect such edges without too much decreasing the
 4 threshold on `dct_dc_size_luminance` or `dct_dc_size_chrominance` which will increase the
 5 noise in the edge signal, and/or to achieve more precise vertical edge localization
 6 whenever it is required, the magnitude and the sign of the lowest order AC horizontal
 7 frequency coefficient (C_{01}) can be inspected in both the prior block and the current block.
 8 Whether or not there is a vertical edge within the bounds of the prior block or the current
 9 block not aligned with the vertical boundary between these two blocks, should be
 10 indicated by whether or not there is a lowest order AC horizontal frequency coefficient
 11 (C_{01}) with considerable magnitude in either of the two blocks, and if there is a lowest
 12 order AC horizontal frequency coefficient (C_{01}) with considerable magnitude in either of
 13 the two blocks, then the sign of this lowest order AC horizontal frequency coefficient
 14 (C_{01}) should indicate whether the gradient of the vertical edge is positive or negative in
 15 the horizontal direction.

16 FIG. 24 shows logic responsive to attributes of only a current block for detecting
 17 when within the current block there is (more likely than not) an almost vertical edge with
 18 a negative horizontal gradient component and when within the current block there is
 19 (more likely than not) an almost vertical edge with a positive horizontal gradient
 20 component. In this case, the MPEG encoded bit-stream segment for the current block is
 21 parsed to determine whether there is a variable-length code for the lowest order AC
 22 horizontal frequency coefficient C_{01} and if so, a signal (`MAG_DCT_01.GT.TC01`) is
 23 provided indicating whether or not the magnitude of the coefficient C_{01} is greater than a

1 threshold T_{C01} , and a signal (SIGN_DCT_ C_{01}) is provided indicating the sign of the
 2 coefficient C_{01} . The threshold T_{C01} can be calculated by taking the inner product of the
 3 DCT basis image associated with the coefficient C_{01} , with an ideal vertical step edge of
 4 the minimal contrast level desired to be detected. Furthermore, this threshold value can
 5 be replaced by a fraction of it to account for the influences of non-ideal (finite gradient)
 6 edges which are neither exactly vertical nor centered within the current block. An OR
 7 gate 190 combines the value of the EDGE_SIGNAL for the current block with the value
 8 of $MAG_DCT_C_{01} > T_{C01}$ for the current block, to enable edge detection even when edge
 9 detection based on dct_dc_size_luminance or dct_dc_size_chrominance thresholding
 10 fails (EDGE_SIGNAL equals logic 0) due to an almost vertical edge in the current block
 11 translated significantly away from the prior block. An exclusive-OR gate 191 compares
 12 the sign of the DCT DC coefficient C_{00} of the current block with the sign of the DCT AC
 13 coefficient C_{01} of the current block to produce a logic 1 if the signs are different and a
 14 logic 0 if the signs are the same. Having opposite signs for the coefficients C_{00} and C_{01} is
 15 consistent with having a simple almost vertical step edge within the current block when
 16 for the current block PD_DCT_ C_{00} equals 0 and having the same sign for the coefficients
 17 C_{00} and C_{01} is consistent with having a simple almost vertical step edge within the current
 18 block when for the current block PD_DCT_ C_{00} equals 1 and all other cases in general
 19 imply the existence of a more complex pattern in that image region possibly with
 20 multiple edges. An exclusive-OR gate 195 combines the output of the exclusive-OR gate
 21 191 with the logical signal PD_DCT_ C_{00} for the current block to signal the potential
 22 presence of a simple almost vertical step edge within the current block. An AND gate
 23 192 produces a signal that is a logic 1 indicating that more likely than not there is an

1 almost vertical edge in the current block with a positive horizontal gradient component if
 2 the edge signal or $MAG_DCT_C_{01}.GT.T_{C01}$ is a logic 1, the sign of the lowest order DCT
 3 AC horizontal frequency coefficient C_{01} is negative ($SIGN_DCT_C_{01}$ is a logic 1) for the
 4 current block, and the exclusive-OR gate 195 produces a logic 1 indicating that either the
 5 sign of the coefficient C_{00} for the current block is different from the sign of the
 6 coefficient C_{01} for the current block and $PD_DCT_C_{00} = 0$ or the sign of the coefficient
 7 C_{00} for the current block is the same as the sign of the coefficient C_{01} for the current
 8 block and $PD_DCT_C_{00} = 1$. An inverter 193 inverts the value of $SIGN_DCT_C_{01}$ for
 9 the current block to produce a logic 1 when the sign of the coefficient C_{01} is positive for
 10 the current block. An AND gate 194 produces a signal that is a logic 1 indicating that
 11 more likely than not there is an almost vertical edge in the current block with a negative
 12 horizontal gradient component if the edge signal or $MAG_DCT_C_{01}.GT.T_{C01}$ is a logic 1,
 13 the sign of the lowest order DCT AC horizontal frequency coefficient C_{01} is positive
 14 ($SIGN_DCT_C_{01}$ is a logic 0) for the current block, and the exclusive-OR gate 195
 15 produces a logic 1 indicating that either the sign of the coefficient C_{00} for the current
 16 block is different from the sign of the coefficient C_{01} for the current block and
 17 $PD_DCT_C_{00} = 0$ or the sign of the coefficient C_{00} for the current block is the same as the
 18 sign of the coefficient C_{01} for the current block and $PD_DCT_C_{00} = 1$.

19 FIG. 25 shows logic responsive to attributes of only a current block for detecting
 20 when within the current block there is (more likely than not) an almost horizontal edge
 21 with a negative vertical gradient component and when within the current block there is
 22 (more likely than not) an almost horizontal edge with a positive vertical gradient
 23 component. For this logic to be applicable, in the adopted block scan order, either the

1 current block should be located in the block row above the previous block or the current
2 block should be located in the block row below the previous block. The MPEG encoded
3 bit-stream segment for the current block is parsed to determine whether there is a
4 variable-length code for the lowest order AC vertical frequency coefficient C_{10} and if so,
5 a signal ($MAG_DCT_C_{10}.GT.T_{C10}$) is provided indicating if or not the magnitude of the
6 coefficient C_{10} is greater than a threshold T_{C10} , and a signal ($SIGN_DCT_C_{10}$) is provided
7 indicating the sign of the coefficient C_{10} . The threshold T_{C10} can be calculated by taking
8 the inner product of the DCT basis image associated with the coefficient C_{10} , with an
9 ideal horizontal step edge of the minimal contrast level desired to be detected.
10 Furthermore, this threshold value can be replaced by a fraction of it to account for the
11 influences of non-ideal (finite gradient) edges which are neither exactly horizontal nor
12 centered within the current block. An OR gate 200 combines the value of the
13 $EDGE_SIGNAL$ for the current block with the value of $MAG_DCT_C_{10}.GT.T_{C10}$ for the
14 current block, to enable edge detection even when edge detection based on
15 $dct_dc_size_luminance$ or $dct_dc_size_chrominance$ thresholding fails ($EDGE_SIGNAL$
16 equals logic 0) due to an almost horizontal edge in the current block translated
17 significantly away from the prior block. An exclusive-OR gate 201 compares the sign of
18 the DCT DC coefficient C_{00} of the current block with the sign of the DCT AC coefficient
19 C_{10} of the current block to produce a logic 1 if the signs are different and a logic 0 if the
20 signs are the same. Having opposite signs for the coefficients C_{00} and C_{10} is consistent
21 with having a simple almost horizontal step edge within the current block when for the
22 current block $PD_DCT_C_{00}$ equals 1 and having the same sign for the coefficients C_{00}
23 and C_{10} is consistent with having a simple almost horizontal step edge within the current

1 block when for the current block $PD_DCT_C_{00}$ equals 0 and all other cases in general
2 imply the existence of a more complex pattern in that image region possibly involving
3 multiple edges. An exclusive-OR gate 205 combines the output of the exclusive-OR gate
4 201 with the negation of the logical signal $PD_DCT_C_{00}$ for the current block provided
5 by the inverter 206 to signal the potential presence of a simple almost horizontal step
6 edge within the current block. An AND gate 202 produces a signal that is a logic 1
7 indicating that more likely than not there is an almost horizontal edge in the current block
8 with a positive vertical gradient component if the edge signal or
9 $MAG_DCT_C_{10}.GT.T_{C10}$ is a logic 1, the sign of the lowest order DCT AC vertical
10 frequency coefficient C_{10} is negative ($SIGN_DCT_C_{10}$ is a logic 1) for the current block,
11 and the exclusive-OR gate 205 produces a logic 1 indicating that either the sign of the
12 coefficient C_{00} for the current block is different from the sign of the coefficient C_{10} for
13 the current block and $PD_DCT_C_{00} = 1$ or the sign of the coefficient C_{00} for the current
14 block is the same as the sign of the coefficient C_{10} for the current block and
15 $PD_DCT_C_{00} = 0$. An inverter 203 inverts the value of $SIGN_DCT_C_{10}$ for the current
16 block to produce a logic 1 when the sign of the coefficient C_{10} is positive for the current
17 block. An AND gate 204 produces a signal that is a logic 1 indicating that more likely
18 than not there is an almost horizontal edge in the current block with a negative vertical
19 gradient component if the edge signal or $MAG_DCT_C_{10}.GT.T_{C10}$ is a logic 1, the sign of
20 the lowest order DCT AC vertical frequency coefficient C_{10} is positive ($SIGN_DCT_C_{10}$
21 is a logic 0) for the current block, and the exclusive-OR gate 205 produces a logic 1
22 indicating that either the sign of the coefficient C_{00} for the current block is different from
23 the sign of the coefficient C_{10} for the current block and $PD_DCT_C_{00} = 1$ or the sign of

1 the coefficient C_{00} for the current block is the same as the sign of the coefficient C_{10} for
 2 the current block and $PD_DCT_C_{00} = 0$.

3 As described above, there are various ways of producing edge indicating bits from
 4 which can be calculated auto-coincidence counts and cross-coincidence counts for
 5 detecting scene changes. Another way of extracting features from the edge indicating
 6 data is to trace the edges to detect joints where the edges intersect in order to construct
 7 graphs representing the morphology of objects in a scene. The graph of a current frame
 8 can be compared to a graph of a prior frame to detect significant differences in the
 9 location and number of the joints and in the number of edges that intersect at the joints.
 10 The information about the edge graphs for the I-frames could be stored as I-frame
 11 metadata for use by a media database search engine to find the I-frames that match a
 12 given frame or a search specification in terms of graph characteristics. As such, the edge
 13 graph information can be used for object identification in the compressed data domain.

14 When tracing edges to find joints, the thinning filter of FIG. 9 could be used to
 15 retain less significant edge segments that would provide connections between the joints.
 16 The code length and sign of the DCT DC coefficients could be saved in a frame buffer
 17 for ease of access and comparison when tracing the edges in a vertical, horizontal or
 18 diagonal direction when searching for the joints. It may also be desirable to have more
 19 information about the direction of the edges. One way of extracting more information
 20 about the direction of the edges is to compute an estimate of the gradient vector of an
 21 edge from the lowest AC horizontal frequency DCT coefficient C_{01} and the lowest AC
 22 vertical frequency DCT coefficient C_{10} for the current block.

23 FIG. 26 is a flow chart of a procedure for estimating the gradient vector of an

1 edge. In a first step 211, the processor decodes an "x" component G_{cx} of the gradient
 2 vector for the current block as the value of the lowest AC horizontal frequency DCT
 3 coefficient C_{01} for the current block with a sign inversion. In step 212, the processor
 4 decodes a "y" component G_{cy} of the gradient vector for the current block as the value of
 5 the lowest AC vertical frequency DCT coefficient C_{10} for the current block with a sign
 6 inversion. Both sign inversions are necessary owing to the reference polarities of the
 7 basis images associated with the DCT coefficients C_{01} and C_{10} as well as the standard
 8 positive senses' being fixed as towards right and towards down respectively along the
 9 horizontal and vertical directions of an image. In step 213, the gradient vector of the
 10 edge is estimated as $G_{cx}\mathbf{x} + G_{cy}\mathbf{y}$.

11 In view of the above, there have been described very fast and computationally
 12 efficient methods of edge detection for block coded video and scene change detection for
 13 MPEG video. Edges are detected in block coded video by a threshold comparison upon
 14 the lengths of variable-length codes used for encoding the differential DC coefficients of
 15 the pixel blocks. A thinning filter compares the code lengths of the differential DC
 16 coefficients of adjacent blocks in order to retain the edge indications of more significant
 17 edges and to exclude the edge indications of less significant edges. The edge indications
 18 can be split into substantially independent channels for luminance or chrominance, for
 19 edges having positive or negative horizontal gradient components as indicated by the sign
 20 of the differential DC coefficient, the prediction direction of the differential DC
 21 coefficient, and the sign of the lowest horizontal frequency DCT coefficient for the pixel
 22 block containing the edge or for edges having positive or negative vertical gradient
 23 components as indicated by the sign of the differential DC coefficient, the prediction

1 direction of the differential DC coefficient, and the sign of the lowest vertical frequency
2 DCT coefficient for the pixel block containing the edge. An approximate gradient vector
3 of an edge can be computed from the lowest horizontal frequency DCT coefficient and
4 the lowest vertical frequency DCT coefficient. The edge indications for successive
5 frames in an MPEG sequence can be compared to each other in various ways in order to
6 detect scene changes.

7